

## CLAIMS

What is claimed is:

- 1 1. A method comprising:
  - 2 creating a guest translation data structure to be used by a guest operating
  - 3 system for address translation operations;
  - 4 creating an active translation data structure based on the guest translation data
  - 5 structure; and
  - 6 periodically modifying content of the active translation data structure to
  - 7 conform to content of the guest translation data structure, the content of the active
  - 8 translation data structure being used by a processor to cache address translations in a
  - 9 translation-lookaside buffer (TLB).
- 1 2. The method of claim 1 further comprising emulating functionality of the TLB
  - 2 in response to an address-translation operation performed by the guest operating
  - 3 system.
- 1 3. The method of claim 1 further comprising:
  - 2 receiving control over an event initiated by guest software; and
  - 3 evaluating the event to identify a cause of the event.
- 1 4. The method of claim 3 further comprising determining that the event is caused
  - 2 by an inconsistency between the content of the active translation data structure and
  - 3 the content of the guest translation data structure.

1     5.     The method of claim 4 further comprising modifying one or more entries in  
2     the active translation data structure that are associated with the event and do not  
3     match corresponding entries in the guest translation data structure.

1        6.        The method of claim 4 further comprising;  
2                comparing the content of the active translation data structure with the content  
3        of the guest translation data structure; and  
4                modifying all entries in the active translation data structure that do not match  
5        corresponding entries in the guest translation data structure.

1        7.        The method of claim 3 further comprising:  
2                determining that the event is associated with a page fault that would occur  
3        under normal operation of the guest software; and  
4                passing control over the event to the guest software.

1        8.        The method of claim 3 further comprising:  
2                determining that the event indicates an attempt of the guest software to  
3        manipulate the TLB; and  
4                modifying the content of the active translation data structure.

1        9.        The method of claim 3 further comprising:  
2                determining that the event is associated with a page fault that would not occur  
3        under normal operation of the guest software; and  
4                modifying the content of the active translation data structure.

1 10. The method of claim 1 further comprising initializing the active translation  
2 data structure by invalidating every entry in the active translation data structure to  
3 emulate the TLB with no entries.

1 11. The method of claim 1 further comprising maintaining the active translation  
2 data structure for each virtual machine running on a computer.

1 12. The method of claim 1 further comprising maintaining the active translation  
2 data structure for each set of privilege levels that can be distinguished by page-based  
3 protection.

1 13. The method of claim 3 further comprising:  
2 determining that the event is caused by an attempt of the guest software to  
3 change a privilege level; and  
4 determining whether the changed privilege level may affect page-based  
5 protection.

1 14. An apparatus comprising:  
2 a guest translation data structure to translate virtual memory addresses into  
3 physical memory addresses, the guest translation data structure being managed by a  
4 guest operating system;  
5 an active translation data structure to contain data derived from content of the  
6 guest translation data structure, the active translation data structure being managed by  
7 a virtual machine monitor (VMM); and

8 a translation-lookaside buffer (TLB) to store address translations derived from  
9 the active translation data structure, the TLB being managed by a processor.

1 15. The apparatus of claim 14 further comprising the VMM to periodically modify  
2 content of the active translation data structure to conform to content of the guest  
3 translation data structure and to emulate functionality of the TLB in response to an  
4 address translation operation performed by the guest operating system.

1 16. The apparatus of claim 15 wherein the VMM is to receive control over an  
2 event initiated by guest software and to evaluate the event to identify a cause of the  
3 event.

1 17. The apparatus of claim 16 wherein the VMM is to further determine that the  
2 event is caused by an inconsistency between the content of the active translation data  
3 structure and the content of the guest translation data structure.

1 18. The apparatus of claim 17 wherein the VMM is to further modify one or more  
2 entries in the active translation data structure that are associated with the event and do  
3 not match corresponding entries in the guest translation data structure.

1 19. The apparatus of claim 17 wherein the VMM is to further compare the content  
2 of the active translation data structure with the content of the guest translation data  
3 structure and to modify all entries in the active translation data structure that do not  
4 match corresponding entries in the guest translation data structure.

1 20. The apparatus of claim 16 wherein the event indicates an attempt of the guest  
2 software to manipulate the TLB.

1 21. The apparatus of claim 16 wherein the event is associated with a page fault  
2 generated by a processor in response to an operation performed by the guest software.

1 22. The apparatus of claim 14 wherein the active translation data structure is  
2 maintained for each virtual machine running on a computer.

1 23. The apparatus of claim 14 wherein the active translation data structure is  
2 maintained for each set of privilege levels that can be distinguished by page-based  
3 protection.

1 24. A system comprising:  
2 a memory to store a guest translation data structure used by a guest operating  
3 system for address translation operations and an active translation data structure  
4 derived content from content of the guest translation data structure and managed by a  
5 virtual machine monitor (VMM); and  
6 a processor, coupled to the memory, to contain a translation-lookaside buffer  
7 (TLB), to cache address translations derived from the content of the active translation  
8 data structure in the TLB, and to manage the TLB.

1 25. The system of claim 24 wherein the processor is to transfer control over a  
2 particular event initiated by guest software to the VMM.

1 26. The system of claim 24 wherein the memory is to store the active translation  
2 data structure for each virtual machine.

1 27. The system of claim 24 wherein the memory is to store the active translation  
2 data structure for each set of privilege levels that can be distinguished by page-based  
3 protection.

1 28. A computer readable medium that provides instructions, which when executed  
2 on a processor, cause said processor to perform operations comprising:  
3 creating a guest translation data structure to be used by a guest operating  
4 system for address translation operations;  
5 creating an active translation data structure based on the guest translation data  
6 structure; and  
7 periodically modifying content of the active translation data structure to  
8 conform to content of the guest translation data structure, the content of the active  
9 translation data structure being used by a processor to cache address translations in a  
10 translation-lookaside buffer (TLB).

1 29. The computer readable medium of claim 28 comprising further instructions  
2 causing the processor to perform operations comprising:  
3 receiving control over an event initiated by guest software;  
4 determining that the event is caused by an inconsistency between the content  
5 of the active translation data structure and the content of the guest translation data  
6 structure; and

09752433-122700  
7           modifying one or more entries in the active translation data structure that are  
8   associated with the event and do not match corresponding entries in the guest  
9   translation data structure.

1   30.    The computer readable medium of claim 28 comprising further instructions  
2   causing the processor to perform operations comprising:  
3           receiving control over an event initiated by guest software;  
4           determining that the event is caused by an inconsistency between the content  
5   of the active translation data structure and the content of the guest translation data  
6   structure;        comparing the content of the active translation data structure with the  
7   content of the guest translation data structure; and  
8           modifying all entries in the active translation data structure that do not match  
9   corresponding entries in the guest translation data structure.